REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
29-11-2010	Workshop Technical Report	Aug 2010 - Nov 2010

4. TITLE AND SUBTITLE

Strategic Directions in Software at Scale

5a. CONTRACT NUMBER
DDRE-ISCS-2010-1

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
0603781D8Z

6. AUTHOR(S)

May, Michael J. (DDR&E/ISCS)
Lee, Edward A. (University of California, Berkeley/CHESS)
Jones, Lindsay E. (DDR&E/ISCS)

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER
N/A

5f. WORK UNIT NUMBER
N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Office of Information Systems & Cyber Security (ISCS)
Office of the Director, Defense Research & Engineering (DDR&E)
1777 N. Kent St., Suite 9030
Rosslyn, VA 22209

The University of California, Berkeley
Center for Hybrid and Embedded Software Systems (CHESS)
337 Cory Hall, #1770
Berkeley, CA 94720-1770

8. PERFORMING ORGANIZATION REPORT NUMBER

DDRE-ISCS-2010-1

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Information Systems & Cyber Security (ISCS)
Office of the Director, Defense Research & Engineering (DDR&E)
1777 N. Kent St, Suite 9030
Arlington, VA 22209

10. SPONSOR/MONITOR'S ACRONYM(S)

DDR&E/ISCS

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
DDRE-ISCS-2010-1

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved For Public Release; Distribution is Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

In August 2010, the Office of Information Systems and Cyber Security (ISCS) within the Office of the Director, Defense Research and Engineering (DDR&E) sponsored the Strategic Directions in Software at Scale (SaS) Workshop, hosted by the University of California, Berkeley. The goals of the workshop were to: identify new ideas and promising research directions in software engineering and computer science achievable in the short-, mid-, and long-term; identify opportunities for collaboration and engage in rich intellectual exchange of technical ideas; create a foundation for developing a DoD roadmap for SaS; and begin to build a case for increasing DoD investment in software engineering and computer science research to strengthen the DoD's software technology base.

Fifteen invited speakers gave presentations in the areas of software synthesis, robust and continuous behavior, temporal semantics, scalable composition, and software engineering process and methodology. Each speaker advocated a particular technical approach that could be the basis for a "Strategic Direction" in software research. To capture the quality and promise of the technical approaches, attendees were asked to rate each presentation with respect to six evaluation criteria, and a weighted rank analysis revealed the three best technical approaches overall as well as the top three performers for each individual evaluation criterion.

ISCS found the workshop extremely useful and felt that it benefitted software researchers and the Department of Defense by assisting with community coordination and increased awareness.

15. SUBJECT TERMS

software, scale, embedded system, CHESS, information systems, cyber security, software synthesis, robust and continuous behavior, temporal semantics, scalable composition, software engineering process and methodology

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Michael J. May
U	U	U	UU	18	**19b. TELEPHONE NUMBER** *(include area code)* 703-696-8012

CLEARANCE REQUEST FOR PUBLIC RELEASE OF DEPARTMENT OF DEFENSE INFORMATION

(See Instructions on back.)

(This form is to be used in requesting review and clearance of DoD information proposed for public release in accordance with DoDD 5230.9.)

TO: *(See Note)* **Chief, Office of Security Review, 1155 Defense Pentagon, Washington, DC 20301-1155**

Note: Regular mail address shown above. For drop-off/next day delivery, use:
Room 12047, 1777 North Kent Street, Rosslyn, VA 22209-2133

1. DOCUMENT DESCRIPTION

a. TYPE	b. TITLE
Presentation	Strategic Directions in Software at Scale (SaS)
c. PAGE COUNT	d. SUBJECT AREA
15	Information System and Cyber Security

2. AUTHOR/SPEAKER

a. NAME *(Last, First, Middle Initial)*	b. RANK	c. TITLE
Dr. Michael May	Civilian	Associate Director, Software Technologies
d. OFFICE		e. AGENCY
OSD-AT&L/DUSD/RD/ Information Systems		DoD

3. PRESENTATION/PUBLICATION DATA *(Date, Place, Event)*

Held August 18-19, 2010, University of California, Berkeley, Center for the Hybrid and Embedded of Software Systems (CHESS)

4. POINT OF CONTACT

a. NAME *(Last, First, Middle Initial)*	b. TELEPHONE NO. *(Include Area Code)*
May, Michael J.	703-696-8012

5. PRIOR COORDINATION

a. NAME *(Last, First, Middle Initial)*	b. OFFICE/AGENCY	c. TELEPHONE NO. *(Include Area Code)*
	CLEARED For Open Publication NOV 03 2010 5 Office of Security Review Department of Defense	

6. REMARKS

Note: to the PR request that the material was pulled together from other PR sources. If they want to know the owners/generators of the material, the reviewer should look at the acknowledgement slide

Please contact Nika Jackson at 703-588-7444 or Dr. Michael May at 703-696-8012, once the presentation is reviewed and ready for pick-up.

7. RECOMMENDATION OF SUBMITTING OFFICE/AGENCY

a. THE ATTACHED MATERIAL HAS DEPARTMENT/OFFICE/AGENCY APPROVAL FOR PUBLIC RELEASE *(qualifications, if any, are indicated in Remarks section)* AND CLEARANCE FOR OPEN PUBLICATION IS RECOMMENDED UNDER PROVISIONS OF DODD 5320.9. I AM AUTHORIZED TO MAKE THIS RECOMMENDATION FOR RELEASE ON BEHALF OF:

ODUSD(AT&L)/ DDR&E

b. CLEARANCE IS REQUESTED BY 20101231 *(YYYYMMDD).* 11-S-0340

c. NAME *(Last, First, Middle Initial)*	d. TITLE
May, Michael J.	Associate Director, Software Technologies
e. OFFICE	f. AGENCY
ODUSD(AT&)/RD/ Information Systems and Cyber Security	OSD/AT&L/DDR&E
g. SIGNATURE	h. DATE SIGNED *(YYYYMMDD)*
	2010 11 02

DD FORM 1910, JAN 2006 PREVIOUS EDITION MAY BE USED. FormFlow/Adobe Professional 6.0

Strategic Directions

in

Software at Scale (SaS)

Held August 18-19, 2010

Hosted by

The University of California, Berkeley

Center for Hybrid and Embedded Software Systems (CHESS)

In coordination with

The Office of Information Systems and Cyber Security,

Office of the Director, Defense Research and Engineering

Workshop Report

September, 27 2010

CLEARED
For Open Publication
NOV 03 2010
Office of Security Review
Department of Defense

11-S-0340

Strategic Directions

in

Software at Scale (SaS)

Held August 18-19, 2010

Hosted by

The University of California, Berkeley
Center for Hybrid and Embedded Software Systems (CHESS)
Berkeley, CA

In coordination with

The Office of Information Systems and Cyber Security (ISCS)
Office of the Director, Defense Research and Engineering (DDR&E)
Rosslyn, VA

Workshop Technical Report

DDRE-ISCS-2010-1

November 2010

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

Final Workshop Technical Report

DDRE-ISCS-2010-1

STRATEGIC DIRECTIONS IN SOFTWARE AT SCALE

Contract W911NF-07-2-0019

Prepared for

Office of Information Systems & Cyber Security (ISCS)

Office of the Director, Defense Research & Engineering (DDR&E)

Rosslyn, VA

For the period

August 2010 – November 2010

Prepared by

Michael J. May, DDR&E/ISCS

Edward A. Lee, University of California, Berkeley

Lindsay E. Jones, DDR&E/ISCS

The University of California, Berkeley
Center for Hybrid and Embedded Software Systems (CHESS)
Berkeley, CA

In coordination with

The Office of Information Systems and Cyber Security (ISCS)
Office of the Director, Defense Research and Engineering (DDR&E)
Rosslyn, VA

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

Table of Contents

List of Figures and Tables

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

1. EXECUTIVE SUMMARY

In August 2010, the Office of Information Systems and Cyber Security (ISCS) within the Office of the Director, Defense Research and Engineering (DDR&E) sponsored the Strategic Directions in Software at Scale (SaS) Workshop. The SaS Workshop was hosted by the University of California, Berkeley. The goals of the workshop were to:

- Identify new ideas and promising research directions in software engineering and computer science achievable in the short-, mid-, and long-term.
- Identify opportunities for collaboration and engage in rich intellectual exchange of technical ideas.
- Create a foundation for developing a DoD roadmap for SaS.
- Begin to build a case for increasing DoD investment in software engineering and computer science research to strengthen the DoD's software technology base.

Fifteen invited speakers gave presentations in the areas of software synthesis, robust and continuous behavior, temporal semantics, scalable composition, and software engineering process and methodology. Each speaker advocated a particular technical approach that could be the basis for a "Strategic Direction" in future software research. To capture the quality and promise of the technical approaches, attendees were asked to rate each presentation with respect to six evaluation criteria.

The overall best technical approaches, as assessed by the attendees, were "Temporal Semantics in Concurrent and Distributed Software"—Edward Lee, "Is Distributed Consistency Scalable?"—Ken Birman, and "The Effect of Software (and Communication) Reliability and Security on Control Systems"—Bruno Sinopoli.

Table ES-1 depicts the top performers in rank order for each Evaluation Criteria (EC) as determined by the weighted rank analysis described in section 4.2. Workshop presentations have been archived at http://chess.eecs.berkeley.edu/conferences/10/SDISAS/index.htm. Hyperlinks to each individual presentation are included in the Appendix.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

Evaluation Criteria		Technical Approach	Advocate
EC1	How does the goal of the research compare to the state of the art?	Temporal Semantics in Concurrent and Distributed Software	Lee
		From Formal Verification to Synthesis	Alur
		Control Software for Systems that Change Structure	Sengupta
EC2	Is the research unique?	Temporal Semantics in Concurrent and Distributed Software	Lee
		The Effect of Software (and Communication) Reliability and Security on Control Systems	Sinopoli
		Computer Aided Programming: Enabling Software at Scale	Solar-Lezama
EC3	Who would use the knowledge?	Composition at Scale	Sztipanovits
		Temporal Semantics in Concurrent and Distributed Software	Lee
		Opportunity-Centered Software Development Environments	Sullivan
EC4	How much will it cost?	Engineering Processes that Engineer Scalable Systems	Osterweil
		Synthesis of Provably-Correct Software Using Discrete Control Theory	Wang
		Synthesis for Software Security	Foster
EC5	How long will it take?	Engineering Processes that Engineer Scalable Systems	Osterweil
		Synthesis of Provably-Correct Software Using Discrete Control Theory	Wang
		Is Distributed Consistency Scalable?	Birman
EC6	What are the measures of success?	Is Distributed Consistency Scalable?	Birman
		Quantitative Verification and Synthesis of Systems	Seshia
		Control Software for Systems that Change Structure	Sengupta

Table ES-1: Top Three Technical Approaches (in Rank Order) for Each Evaluation Criterion

The ISCS office found the workshop extremely useful and felt that it benefitted software researchers by assisting with community coordination and increased awareness. Several suggestions on how to improve such workshops in the future were also made. They included:

- Define objectives more clearly,
- Push for participation from other Government agencies and industrial organizations,
- Evolve and enhance the evaluation criteria and assessments,
- Enhance workshop structure.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

2. INTRODUCTION

2.1. Workshop Purpose and Goals

Software has become a critical enabler of our nation's defense systems and is rapidly and continually increasing in size, scale, and complexity. The Department of Defense (DoD) as well as the industrial base continues to encounter difficulties in successfully deploying software-intensive systems with desired functionality under cost and schedule constraints. Shortcomings and failure to successfully execute software-intensive systems can often be attributed to underpowered software development technologies which are not capable of addressing the scale, complexity, and capability required of today's systems. These underpowered technologies may be symptomatic of lacking investments in software engineering and computer science research and development (R&D) at the fundamental level, and a decline in DoD software expertise and knowledge assets.

In an effort to begin to confront these issues, the Office of Information Systems and Cyber Security (ISCS) within the Office of the Director, Defense Research and Engineering (DDR&E), in conjunction with the University of California, Berkeley was motivated to bring together a forum of the best thinkers across academia, industry, and Government to advocate and promote ideas with potential to dramatically improve our collective ability to build, evolve, and use large software systems; this resulted in plans for a 2010 Strategic Directions in Software at Scale (SaS) Workshop. The goals in hosting this workshop were to leverage the candidate technical approaches presented and discussions provoked to:

- Identify new ideas and promising research directions in software engineering and computer science achievable in the short-, mid-, and long-term.
- Identify opportunities for collaboration and engage in rich intellectual exchange of technical ideas.
- Create a foundation for developing a DoD roadmap for SaS.
- Begin to build a case for increasing DoD investment in software engineering and computer science research to strengthen the DoD's software technology base.

The remainder of this report summarizes the SaS workshop approach, technology areas explored, and overall results and conclusions about the technical approaches advocated by workshop speakers.

2.2 Workshop Approach

The Strategic Directions in Software at Scale workshop was invitation-only. Workshop participants included researchers, practitioners, and program managers from industry, academia, and Government. Potential workshop attendees were recommended Berkeley researchers and approved by ISCS staff based on expertise and areas of interest, involvement in the research community, and participation in DoD-sponsored software R&D programs. The ISCS staff also augmented the Berkeley-recommended list of attendees to broaden the intellectual base at the workshop.

The following high-level "technical focus areas" were chosen to establish a manageable scope for the workshop discussions (these are discussed further in section 2.3):

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

(1) Software Synthesis
(2) Temporal Semantics
(3) Scalable Composition
(4) Robust and Continuous behavior
(5) Secure Composition
(6) Process and Methodology

A subset of the invited workshop participants was invited to lead presentation/discussion sessions to advocate a technical approach or strategic direction in one of the six technical focus areas. Each session leader was given 20 minutes to advocate and make a case for their specific strategic direction or technical approach. This was followed by 20 minutes of group discussion during which the leader could pose questions to the group to stimulate moderated debate and dialogue.

Workshop speakers were provided with general guidance on how to structure their sessions, including:

- A description and overview of the technical direction advocated.
- The challenges or problems addressed and limitations of current practice.
- Novel technical aspects of the promising approach and evidence to support why it will work.
- Expected payoff including metrics that could assess success.
- Risk factors if the direction is not pursued, and the likelihood of a dead-end.

A representative example was also provided as guidance for formulating arguments and discussion questions:

(20 minutes) Session lead makes a case for the importance of pursuing research in the area of Temporal Semantics, for example:

- Argue that, by choice, computer science has omitted timing from the semantics of programming. The underlying technology, however, is very capable of precise and reliable timing. Argue for potential benefits of integrating timing into the semantics of programs. Risk factors include unknown effects from having to redesign much of the abstraction stack, from instruction set architectures (ISAs) up through operating systems and networks.

(20 minutes) Session lead poses thought-provoking questions and facilitates group discussion:

- If computation and networking speeds continue to improve, can we just circumvent the problem by over provisioning?
- What proportion of software problems arise from uncontrolled or unexpected timing of interaction between software components?
- Are there intermediate solutions that do not require redoing much of what computer science has done for the last 40 years?
- How long might it take for investment in research to lead to payoffs?

The workshop was held over the course of two days with 15 speaker/discussion sessions, as well as an additional separate session for “tweets” during which workshop participants could take five minutes to advocate for a topic, technical or otherwise, related to software research, development, acquisition, etc.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

The following section provides a description of the higher-level technology areas explored as well as the specific technical approaches and strategic directions proposed and discussed over the course of the workshop.

3. SOFTWARE TECHNOLOGY FOCUS AREAS AND TECHNICAL APPROACHES PROPOSED

Advocates were invited to make a case for a strategic direction or technical approaches in one of six software technology focus areas. General descriptions of each of the focus areas are provided in this section, and the technical approaches and strategic directions proposed listed beneath. As expected, several of the technical approaches crossed multiple technology focus areas; technical approaches are listed below based on the focus area with which they most closely align.

(1) Software Synthesis (SS)

The notion is that software implementations can be computed from abstract and incomplete specifications with systematic exploration of the alternative implementations. The goal is to leverage advances in program modeling and analysis to be able to rule out undesirable implementations quickly and guide selection of desirable implementations, and to perform automatic code generation for those implementations.

From Formal Verification to Synthesis
Rajeev Alur, Professor, University of Pennsylvania

Scalable Methods for Managing Uncertainty in System Design
Andrzej Banaszuk, Fellow, United Technologies Research Center

Synthesis for Software Security
Jeffrey Foster, Professor, University of Maryland

Computer Aided Programming: Enabling Software at Scale
Armando Solar-Lezama, Assistant Professor, Massachusetts Institute of Technology

Synthesis of Provably-Correct Software Using Discrete Control Theory
Yin Wang, Research Scientist, HP Labs

(2) Temporal Semantics (TS)

Cyber-physical systems integrate computing and networking with physical processes. The temporal dynamics of software and networks becomes critical to predicting and controlling the interactions of system components. However, nearly all current software abstractions omit time. The theme of this focus area is to investigate the potential impact and technical implications of modifying these abstractions to embrace temporal dynamics.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

Temporal Semantics in Concurrent and Distributed Software
Edward Lee, Professor, University of California, Berkeley

Software at Scale: Temporal Semantics
Vijay Saraswat, Member of Research Staff, IBM

Quantitative Verification and Synthesis of Systems
Sanjit Seshia, Assistant Professor, University of California, Berkeley

The Effect of Software (and Communication) Reliability and Security on Control Systems
Bruno Sinopoli, Assistant Professor, Carnegie Mellon University

(3) Scalable Composition (SC)

Many complex designs fail at system integration because of underspecified interfaces, unstated assumptions, or unexpected interference between components. This focus area addresses the problem through mechanisms for clarifying interfaces of components and ensuring correct composition.

Composition at Scale
Janos Sztipanovits, Professor, Vanderbilt University

(4) Robust and Continuous Behavior (RCB)

Software tends to fail catastrophically, with return to known good state (e.g. rebooting) being a dominant recovery method. This focus area addresses approaches to achieving robust and continuous behaviors, where "continuous" means that small changes have small effects.

Is Distributed Consistency Scalable?
Ken Birman, Professor, Cornell University

Control Software for Systems that Change Structure
Raja Sengupta, Associate Professor, University of California, Berkeley

(5) Secure Composition (SC2)

Complex systems constructed by composing diverse components frequently suffer from interference, where one component disrupts another. This focus area examines mechanisms by which subsystems can be composed with assurances of non-interference. Possible approaches include game-theoretic formulations.

This focus area did not receive any submissions from workshop participants.

(6) Process and Methodology (PM)

Reliable, repeatable production relies on well-understood and well-executed processes. Metrics for assessing quality are required. Further, the production of software at scale requires that the functions of

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

both process and measurement scale efficiently to large or complex systems. This focus area examines such scalable processes and measures related to system components, networks, and human and organizational components.

Reliability and Robustness of Large-Scale Systems
John Goodenough, Fellow, Carnegie Mellon University Software Engineering Institute

Engineering Processes that Engineer Scalable Systems
Lee Osterweil, Professor, University of Massachusetts, Amherst

Opportunity-Centered Software Development Environments
Kevin Sullivan, Professor, University of Virginia and Visiting Scientist, Carnegie Mellon University Software Engineering Institute

A separate "tweet" session allowing participants to quickly argue for a topic related to software at scale included the following:

Real Complexity
Brian Murray, Group Leader, United Technologies Research Center

Leadership Challenges for SaS Development
Edgar Dalrymple, Future Combat Systems Program, US Army Aviation and Missile Research Development and Engineering Center

Virtualization and Isolation
Christoph Kirsch, Postdoctoral Researcher, University of Salzburg

Toward a Science of Software Development
David Luginbuhl, Mathematics, Information and Life Sciences Directorate, Air Force Office of Scientific Research (for Jim Kirby, Center for High Assurance Computer Systems, Naval Research Laboratory)

Problems in Cyber Security
Glenn Racine, Network Sciences Division, Computational and Information Sciences Directorate, Army Research Laboratory

Complex Systems
Edward Lee, Professor, University of California, Berkeley

A workshop agenda is included in the Appendix of this document. This agenda includes hyperlinks to the detailed presentations for each of the technical approaches listed above.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

4. ASSESSMENT OF TECHNICAL APPROACHES

4.1 Evaluation Criteria for Assessment of Technical Approaches

Six evaluation criteria (ECs) were developed as a means to enable "standardized" assessments of each of the technical approaches presented during the workshop to gain a sense from the group about research priorities, risk factors, and the promise of the various technical approaches. These ECs are described in Table 1. Sample responses were provided to guide workshop participants in addressing each of the ECs on a consistent scale to allow for subsequent data analysis. Participants were instructed to select one response per EC among "Best Case," "Middle Case," and "Worst Case" for each technical approach presented. Participants were also encouraged to provide written commentary to supplement their responses or reinterpret the evaluation criteria in free space provided.

Evaluation Criteria (EC)	Sample Response	
EC1: How does the goal of the research compare to the current state of the art?	Provides revolutionary understanding; success would be a breakthrough	Best Case
	Provides new knowledge	Middle Case
	Little or no apparent knowledge gain	Worst Case
EC2: Is the research unique?	Truly novel approach; trailblazing	Best Case
	Extends known concepts in novel ways	Middle Case
	Marginally different from previous work	Worst Case
EC3: Who would use the knowledge?	Obvious universal applications; like GUIs over command-line	Best Case
	Solid, but limited user-base: industry, military, academia	Middle Case
	Applications are unclear or focused on very small communities	Worst Case
EC4: How much will it cost?	Minimal investment required: mainly theoretical investigation, a few good minds.	Best Case
	Significant brain power. Plus test tools, computing hours, and lab space.	Middle Case
	Major research infrastructure: specialized hardware and software, specialized test ware	Worst Case

(Continued on next page)

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

EC5: How long will it take?	I can write the abstract now for a Spring conference.	Best Case
	There known questions which have to be answered. 2 -3 years perhaps.	Middle Case
	Surely, new questions will arise. No confident time frame can be established.	Worst Case
EC6: What are the measures of success?	Established frameworks already exist to measure success in this area.	Best Case
	The measures of "overall effectiveness" are mostly known, but diagnostic and detailed performance metrics are not fully-established.	Middle Case
	It isn't clear how the research goal's effect on anything could be evaluated.	Worst Case

Table 1: Workshop Evaluation Criteria

While the assessment scale is admittedly imperfect, and there is a certain level of interpretation influencing responses, it did provide a manageable means for quantifying and analyzing the reactions of workshop participants for each of the technical approaches proposed. The anonymity of responses presumably enabled attendees to be more frank than they might be speaking openly or if polled in real-time during the workshop. Many participants did take the opportunity to expand on their responses and offer detailed commentary. This data is being used internally by ISCS to interpret and expand on these results.

4.2 Data Analysis Methodology

Raw data responses from the workshop participants' assessments were compiled for each technical approach presented. The number of "Best Case," "Middle Case," and "Worst Case" responses for each EC were divided over the total number of responses for that EC, providing a percent "Best Case" (%BC), percent "Middle Case" (%MC), and percent "Worst Case" (%WC) for each EC. This normalization was necessary due to differences in the number of attendees responding for each technical approach and each EC. A weight of 10 was applied to "Best Case," 5 to "Middle Case," and 1 to "Worst Case," and a weighted sum was calculated to represent a technical approach's performance for each EC, with the highest score possible score being 10, and the lowest 1 for any one EC. As an example, for a representative technical approach X, the weighted sum for EC1 was calculated by:

$$EC1_x = \%BC*10 + \%MC*5 + \%WC*1 \quad (1)$$

To represent each technical approach's overall performance across all ECs, an overall weighted sum was calculated by adding together the weighted sums for $EC1_x$ through $EC6_x$, with each EC having an equal weight:

$$EC_x = EC1_x + EC2_x + EC3_x + EC4_x + EC5_x + EC6_x \quad (2)$$

The weighted sums for each technical approach were rank-sorted from highest to lowest to identify the top three performing technical approaches for each EC and an overall score across ECs.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

The variance for each technical approach was calculated to identify those which varied substantially across the six ECs.

Lastly, by considering each weighted quantity, ECn_x, a component of a six-dimensional vector, **ECx**, unit vectors ($\hat{\mathbf{EC}}_x$) were calculated for each technical approach by dividing by vector magnitudes ($|\mathbf{EC}_x|$). The dot product of $\hat{\mathbf{EC}}_x$ and $\hat{\mathbf{EC}}_y$ was calculated to determine the quantity "cos Θ_{xy}" which can be thought of as the cosine of a generalized angle. This analysis interprets cos Θ as one measure of how similar two sets of EC ratings were. Clearly, two technical approaches receiving the exact same EC ratings would have cos Θ = 1. This quantity was used as a screening tool to look more closely at technical approaches who had several outlying cos Θ values (defined as less than 0.94) when compared to the other approaches.

$$|\mathbf{ECx}| = (\mathbf{EC}_x^{\,2})^{1/2} \quad (3)$$

$$\hat{\mathbf{EC}}_x = \mathbf{EC}_x / |\mathbf{EC}_x| \quad (4)$$

$$\hat{\mathbf{EC}}_x \bullet \hat{\mathbf{EC}}_y = \cos \Theta_{xy} \quad (5)$$

The results of the calculations described in this section are discussed in section 5.1 along with observations and interpretations.

5. RESULTS AND DISCUSSION

5.1 Data Analysis and Results

5.1.1 Analysis of Performance for Each Evaluation Criterion

Results

Table 2 depicts the top performers in rank order for each EC as determined by the weighted rank analysis described in section 4.2.

Evaluation Criteria		Technical Approach	Advocate	Technical Focus Area
EC1	How does the goal of the research compare to the state of the art?	Temporal Semantics in Concurrent and Distributed Software	Lee	TS
		From Formal Verification to Synthesis	Alur	SS
		Control Software for Systems that Change Structure	Sengupta	RCB
EC2	Is the research unique?	Temporal Semantics in Concurrent and Distributed Software	Lee	TS
		The Effect of Software (and Communication) Reliability and Security on Control Systems	Sinopoli	TS
		Computer Aided Programming: Enabling Software at Scale	Solar-Lezama	SS

(Continued on next page)

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

EC3	Who would use the knowledge?	Composition at Scale	Sztipanovits	SC
		Temporal Semantics in Concurrent and Distributed Software	Lee	TS
		Opportunity-Centered Software Development Environments	Sullivan	PM
EC4	How much will it cost?	Engineering Processes that Engineer Scalable Systems	Osterweil	PM
		Synthesis of Provably-Correct Software Using Discrete Control Theory	Wang	SS
		Synthesis for Software Security	Foster	SS
EC5	How long will it take?	Engineering Processes that Engineer Scalable Systems	Osterweil	PM
		Synthesis of Provably-Correct Software Using Discrete Control Theory	Wang	SS
		Is Distributed Consistency Scalable?	Birman	RCB
EC6	What are the measures of success?	Is Distributed Consistency Scalable?	Birman	RCB
		Quantitative Verification and Synthesis of Systems	Seshia	TS
		Control Software for Systems that Change Structure	Sengupta	RCB

Table 2: Top Three Technical Approaches (in Rank Order) for Each Evaluation Criterion

Observations

As shown in Table 2, Edward Lee, Ken Birman, Yin Wang, Lee Osterweil, and Raja Sengupta appear more than once in the top three performers. However, a large number of different talks/advocates appeared in the top three performers across all of the evaluation criteria.

5.1.2 Analysis of Overall Performance

Results

Table 3 depicts the top performers overall as determined by the weighted rank analysis described in section 4.2.

Technical Approach	Advocate	Technical Focus Area
Temporal Semantics in Concurrent and Distributed Software	Edward Lee	TS
Is Distributed Consistency Scalable?	Ken Birman	RCB
The Effect of Software (and Communication) Reliability and Security on Control Systems	Bruno Sinopoli	TS

Table 3: Top Three Technical Approaches (in Rank Order) Overall

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

5.1.3 Variance and Notably Different Evaluation Criteria Patterns

Variance

Three of the technical approaches in particular showed much higher variance across the ECs than the rest: "From Formal Verification to Synthesis" (Alur), "Composition at Scale" (Sztipanovits), and "Temporal Semantics in Concurrent and Distributed Software" (Lee). Lee, the top performer overall, scored extremely high for ECs 1-3 which relate to the technical content, high for EC6 (measures of success), and much lower for ECs 4 and 5 (cost and schedule). Alur scored very high for EC1, high for EC3, and much lower for ECs 4 and 5. Sztipanovits scored extremely high for EC3, high for EC1, and much lower for ECs 2, 4, and 6. Interestingly, all of these approaches received a top three mark in one of ECs 1-3, and scored on the very low end of ECs 4 and 5; this could be evidence of an assumption offered by a workshop participant regarding correlations between the uniqueness and novelty of research and the time and cost required to achieve it. This assertion is discussed further in section 4.1.4.

Dissimilarities when Compared to Other Approaches

Those technical approaches with the largest number of outlying cos Θ_{xy} values were "Engineering Processes that Engineer Scalable Systems" (Osterweil) with five, "Temporal Semantics in Concurrent and Distributed Software" (Lee) with three, and "Composition at Scale" (Sztipanovits) with three. Osterweil performed very high for ECs 4 and 5, slightly lower for ECs 3 and 6, and very low for ECs 1 and 2. As compared to Lee's and Sztipanovits' approaches, described above, Osterweil's approach appears to have been perceived by the attendees to be somewhat orthogonal.

Upon examination, Lee and Sztipanovits' approaches did not score conspicuously diffferently than the others, with the exception of scoring very high in certain ECs (1 and 3 respectively). However, a review of the approaches that were dissimilar to Sztipanovits' approach revealed an interesting feature. The approach "The Effect of Software (and Communication) Reliability and Security on Control Systems" (Sinopoli) was rated very high for EC2 (uniqueness), but lower for EC1 (state of the art) and EC3 (who would use it). It was the only approach to be rated with this pattern.

5.1.4 Additional Observations

Comments on the Quality of Research Presented

As can be inferred from Figure 1, the overall quality of the technical approaches presented was very high. No one EC appeared to dominate the others significantly in terms of performance.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

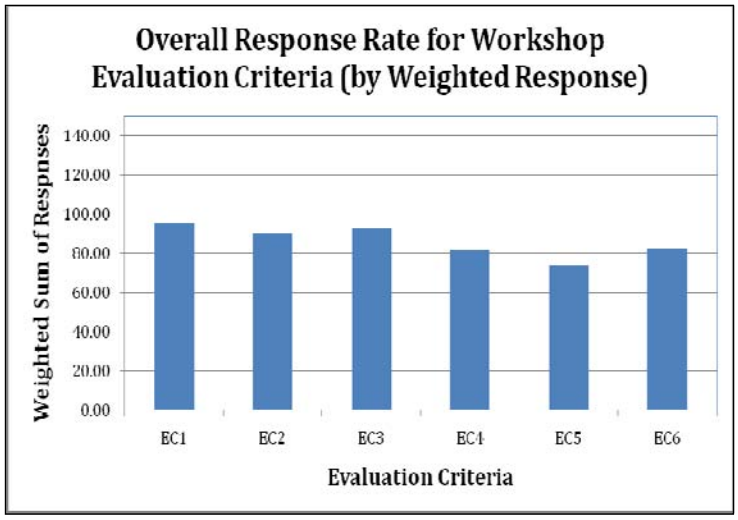

Figure 1: Overall Quality/Performance of Technical Approaches Presented

Correlations Between Evaluation Criteria

Many of the technical approaches which received high marks for comparison to state of the art (EC1) and uniqueness (EC2) received lower marks for cost (EC4) and time to achieve (EC5). While, in many cases, lower scores for ECs 4 and 5 dragged down overall weighted scores, the observation that that ECs 1-2 and 4-5 might be inversely correlated supports the assertion that progressive, unique research inherently lends itself to higher costs and longer time frames, but that perhaps these are the novel, unique challenges that we need to begin addressing now. One might observe that ECs 1-3 seem to come from more of a "research" perspective and ECs 4-6 from more of an "operational" perspective; as such, it is logical to conclude that correlations exist across ECs, and that, depending on interpretations, "Best Case" from an operational perspective might agree with "Worst Case" from a research perspective.

5.2 Workshop Conclusions and Considerations for the Future

We feel that the overall quality of the research presented and technical exchange conducted was extremely high. We hope that attendees perceived value in their participation and are looking forward to future workshops.

We feel that, in addition to the richness of the technical exchange, several benefits were derived from the SaS workshop:

> *Community coordination and community building:* As evidenced by the variety of technical approaches presented and the diversity of backgrounds and research interests, we feel that the SaS workshop took advantage of the opportunity to blend cultural approaches from different, but related, research communities (software, process, control systems). Further, we believe that the technical exchange was made more valuable by the blending of perspectives across members of academia, industry, and the Government. Innovation and progress are often most significant at the intersection of disciplines and outlooks.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

Increased Awareness: We now have a larger pool of researchers from which to draw promising ideas – ideas which have been assessed by peers and experts.

As this was the first of several workshops and activities that we hope to conduct over the next couple of years, we are eager to ensure value, progress, and meaningful return on participants' and sponsors' investment of time and money. As such, we describe below several considerations for future forums:

Clearly define what we intend to do. Ensure that prior to and at the beginning of workshops, goals, intended outcomes, opportunities as a result of participation, and desired input from workshop attendees is *clearly* defined.

Push for participation from other Government agencies and industrial organizations. It is important to foster interagency cooperation to ensure that all avenues are working together, and that researchers and investors are aware of existing opportunities and high-potential ideas.

Evolve and enhance the evaluation criteria and assessments. The six ECs considered at this workshop were established as an "experimental" system to determine if assessments of technical approaches could be structured to facilitate straightforward data analysis.

Continue to enhance workshop structure: Though the two 20 minutes sessions were not strictly adhered to, orchestrating the sessions in this way allowed for lively debate and discussion both during and after the speaker's presentations and ensured ample time for questions and comments. The organizers feel that this structure enabled participants to more thoroughly engage and more comprehensively evaluate the technical approaches presented. We feel it was conducive to rich technical exchange.

In addition, a participant suggested that the organizers consider framing workshops around uses desired by practitioners and end-users, critical problems that need addressing, and capabilities that need to be satisfied. These should be articulated by the DoD.

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited

APPENDIX: WORKSHOP AGENDA

Day 1: Wednesday, 18 August 2010

Time	Agenda Items/Technical Approaches	Speaker
8:30 to 9:45	Workshop Organization	Edward Lee, Berkeley
8:45 to 9:00	Workshop Goals and Overview	Michael May, DDR&E/ISCS, OSD
9:00 to 9:40	*Workshop Introduction – Software at Scale: Critical Defense Issues*	Rich Turner, Stevens
9:40 to 10:20	*From Formal Verification to Synthesis*	*Rajeev Alur, Penn*
10:20 to 10:40	Break	All
10:40 to 11:20	Catch Up	All
11:20 to 12:00	*Is Distributed Consistency Scalable?*	*Ken Birman, Cornell*
12:00 to 1:00	Working Lunch	All
1:00 to 1:40	*Software at Scale: Temporal Semantics*	*Vijay Saraswat, IBM*
1:40 to 2:20	*Synthesis of Provably-Correct Software Using Discrete Control Theory*	*Yin Wang, HP Labs*
2:20 to 3:00	*Quantitative Verification and Synthesis of Systems*	*Sanjit Seshia, Berkeley*
3:00 to 3:30	Break	All
3:30 to 4:10	*Control Software for Systems that Change Structure*	*Raja Sengupta, Berkeley*
4:10 to 5:00	*Synthesis for Software Security*	*Jeffrey Foster, Maryland*
5:00 to 5:30	Wrap up Discussion and Consensus Building	Edward Lee, Berkeley Michael May, DDR&E/ISCS, OSD

Day 2: Thursday, 19 August 2010

Time	Agenda Items/Technical Approaches	Speaker
8:30 to 9:10	*Composition at Scale*	*Janos Sztipanovits, Vanderbilt*
9:10 to 9:50	*Scalable Methods for Managing Uncertainty in System Design*	*Andrzej Banaszuk, UTRC*
9:50 to 10:30	*Engineering Processes that Engineer Scalable Systems*	*Lee Osterweil, Amherst*
10:30 to 10:50	Break	All
10:50 to 11:30	*Opportunity-Centered Software Development Environments*	*Kevin Sullivan, Virginia and SEI*
11:30 to 12:10	*Reliability and Robustness of Large-Scale Systems*	*John Goodenough, SEI, CMU*
12:10 to 1:00	Working Lunch	All
1:00 to 1:40	Tweets (Five Minute Madness)	
1:40 to 2:20	*Computer Aided Programming: Enabling Software at Scale*	*Armando Solar-Lezama, MIT*
2:20 to 3:00	*The Effect of Software (and Communication) Reliability and Security on Control Systems*	*Bruno Sinopoli, CMU*
3:00 to 3:30	Break	All
3:30 to 4:10	*Temporal Semantics in Concurrent and Distributed Software*	*Edward Lee, Berkeley*
5:00 to 5:30	Wrap up Discussion and Consensus Building	Edward Lee, Berkeley Michael May, DDR&E/ISCS, OSD

DISTRIBUTION STATEMENT A: Approved for Public Release; Distribution is Unlimited